

Naturalism Without Mirrors[†]

Huw Price

September 9, 2009

[†]This is a draft table of contents and introductory chapter for a collection of essays to be published by Oxford University Press, New York, with the provisional title *Naturalism Without Mirrors: Pragmatism, Pluralism and the Placement Problem*. Draft copies of the original essays are available online at <http://www.usyd.edu.au/time/price/publications.html>, or via the hyperlinks in the table of contents overleaf.

Contents

1. Introduction: moving the mirror aside.
2. ‘Metaphysical pluralism’, *Journal of Philosophy* 89(1992) 387–409. [[JSTOR](#)]
3. ‘Semantic minimalism and the Frege point’, in Tsohatzidis, S.L.(ed.), *Foundations of Speech Act Theory: Philosophical and Linguistic Perspectives* (Routledge & Kegan Paul, 1994), 132–55. Reprinted with a new postscript in Garrett, B. and Mulligan, K., eds, *Themes From Wittgenstein* (Philosophy Program, RSSH, ANU, 1993), 15–44. [[PDF](#)]
4. ‘Two paths to pragmatism’, in Casati, R. and Tappolet, C., eds., *European Review of Philosophy* 3(1998) 109–47. Originally published in Peter Menzies, ed., *Response-Dependent Concepts*, (Philosophy Program, RSSH, ANU, 1992), 46–82. [[PDF](#)]
5. (With John Hawthorne) ‘How to stand up for non-cognitivists’, *Australasian Journal of Philosophy* 74(1996) 275–292. [[PDF](#)]
6. ‘Naturalism and the fate of the M-worlds’, *Proceedings of the Aristotelian Society*, Supp. Vol. LXXI(1997) 247–267. [[PDF](#)]
7. (With Richard Holton) ‘Ramsey on saying and whistling: a discordant note’, *Noûs* 37:2(2003) 325–341. [[PDF](#)]
8. ‘Truth as convenient friction’, *Journal of Philosophy* 100(2003) 167–190. [[PDF](#)]
9. ‘Naturalism without representationalism’, in David Macarthur and Mario de Caro, eds, *Naturalism in Question* (Harvard University Press, 2004), 71–88. [[PDF](#)]
10. ‘Immodesty without mirrors—making sense of Wittgenstein’s linguistic pluralism’, in Max Kölbel and Bernhard Weiss, eds, *Wittgenstein’s Lasting Significance* (Routledge & Kegan Paul, 2004), 179–205. [[PDF](#)]

11. (With David Macarthur) 'Pragmatism, quasi-realism and the global challenge', in Cheryl Misak, ed., *The New Pragmatists* (Oxford University Press, 2007), 91–120. [\[PDF\]](#)
12. 'The semantic foundations of metaphysics', in Ian Ravenscroft, ed., *Minds, Ethics, and Conditionals: Themes from the Philosophy of Frank Jackson* (Oxford University Press, 2009), 111–140. [\[PDF\]](#)
13. 'Metaphysics after Carnap: the ghost who walks?', in David Chalmers, Ryan Wasserman and David Manley, eds., *Metametaphysics* (Oxford University Press, 2009), 320–346. [\[PDF\]](#)
14. 'One cheer for representationalism?' Forthcoming in Randall Auxier, ed., *The Philosophy of Richard Rorty* (Open Court, Library of Living Philosophers XXXII, 2009). [\[PDF\]](#)

Chapter I

Moving the mirror aside

1.1 The matching game

Imagine a child's puzzle book, arranged like this. The left-hand page contains a large sheet of peel-off stickers, and the right-hand page shows a line drawing of a complex scene. For each sticker – the koala, the boomerang, the Sydney Opera House, and so on – the reader needs to find the unique outline in the drawing with the corresponding shape. The aim of the game is to place all the stickers in their correct locations, in this sense.

Now think of the right-hand page as the world, and the stickers as the collection of all the statements we take to be true of the world. For each such statement, it seems natural to ask what makes it true; what fact in the world has precisely the corresponding “shape”. Within the scope of this simple but intuitive analogy, matching true statements to the world seems a lot like matching stickers to the line drawing.

Moreover, many problems in philosophy seem much like the difficulties the child faces, when some of the stickers are hard to place. In both cases, the difficulties arise from restrictions on the options available on the right-hand side of the game. In the first case, the child has to work within the constraints of the line drawing provided. If she is allowed to draw her own outlines, one for each sticker, the task is bound to be straightforward (engrossing, perhaps, at a certain age, but essentially trivial). In a pre-assigned drawing, however, the required outlines can be concealed or even absent altogether, and hence the puzzle can be difficult or even impossible to complete. Imagine the drawing is constructed from basic shapes like the segments of an orange, for example. We can see how the Opera House fits in, and perhaps the boomerang, but where do we put the koala?

In the philosophical case, similarly, the game is trivial (and not even engrossing, to most temperaments) if for any true statement ‘P’, we are allowed to say that ‘P’ is made

true by the fact that P. It becomes non-trivial when we impose limitations on the facts on the right – restrictions on the available “truthmakers” for the statements on the left.

There are various possible motivations for playing the philosophical version of the game with restrictions of this kind, but let us focus on one in particular – the most influential in contemporary philosophy, I think. It rests on two kinds of intuitions, or theoretical assumptions. The first of these – shared, presumably, with other motivations for the philosophical version of this matching game – is a kind of proto-theory about language, in the light of which the game seems to provide a useful informal model of the relation of language to the world. This proto-theory has to accord a key role to the idea that the function of statements is to “represent” worldly states of affairs, and that true statements succeed in doing so.

It may seem inappropriate to call this assumption a proto-theory. The label “theory” may seem too grand for such an obvious truth, or the label “proto” too tentative for such a well-established canon of philosophy of language. Nothing hangs on the terminology, however. For the moment, the important thing is the role that this assumption – be it trivial truth, proto-theory, or mature canon – plays in giving rise to the most taxing form of the philosophical version of the matching game

If this proto-theory (or whatever) is to be incorporated into a mature scientific theory of the relation of language to the world, then the matching model needs to fit within the scope of a broadly scientific investigation of ourselves, and of the world we inhabit. After all, as we consider the world as scientists, we see ourselves and our language as one small but (to us) rather significant part of it. Hence the second source of the restriction: if the matching model is to be incorporated into a scientific perspective, the perspective itself seems to dictate the shape of the available facts and truthmakers. Roughly, the available shapes are the kinds of outlines recognised by natural science.

Why does this turn out to be a severe constraint, at least *prima facie*? Because there seem to be many true statements that don't line up neatly with any facts of the kind uncovered by natural science. Indeed, the problem cases are not just the classic misfits, such as the (apparent?) truths of aesthetics, morality, and other normative matters, or those of consciousness. Arguably, at least, they include matters much closer to a scientist's heart, such as probability, causation, possibility and necessity, and conditional facts of various kinds; and even, hovering above all, the heavenly truths of mathematics itself.

Thus there is a striking mismatch between the rich world of ordinary discourse and the sparse world apparently described by science. A great deal of work in modern philosophy amounts to attempts to deal with some aspect or other of this mismatch. The project is often called simply naturalism. I shall call it Naturalism, for now, with a capital N, so as to reserve the generic term for a more basic view (with which, as I shall explain, Naturalism itself may well turn out to conflict).

The Naturalist's mantra goes something like this: The only facts there are are the kind of facts recognised by natural science. But it is not this mantra alone which com-

mits Naturalists to their restrictive version of the matching game. In principle, one could endorse the mantra without thinking that the matching game provides a useful model of the relation of language to the world. (Quine provides an example, perhaps, at least under some interpretations.) The puzzle stems from combining the mantra with a piece of proto-science: the kind of proto-theory about language and the world for which the matching game offers a crude model. The proto-theory says that our statements “stand for”, or “represent”, aspects of the world. Big-N Naturalists combine this proto-theory with the mantra’s restriction on the available truthmakers, and it is the combination that leads to the puzzles to which they devote so much philosophical energy.

I have emphasised the role of the proto-theory because it reveals an interesting vulnerability in the Naturalist’s own position. By the Naturalist’s own lights, the proto-theory ought to count as an hypothesis about what it is right to say about language itself, from a naturalistic standpoint. If it turned out to be a bad hypothesis – if better science showed that the proto-theory was a poor theory – then the motivation for the Naturalist’s version of the matching game would be undermined. But it would be undermined from *within* a scientific view of language and its place in the world. In that sense, the undermining wouldn’t be an anti-naturalist conclusion – on the contrary, it would depend on convicting some self-styled naturalists of sub-optimal science.

If we call the proto-theory (big-R) *Representationalism*, then the possibility just mentioned is the possibility that a good naturalistic account of our own linguistic practice might defeat Representationalism – might reveal it to be a poor theory about the relation between language and the world. The result would be naturalism without Representationalism, or naturalism without mirrors.

In ‘Naturalism without Representationalism’ (Ch. 9), I make these points in terms of a distinction between two kinds of naturalism: “object naturalism”, which is the view we’ve just called simply Naturalism, and “subject naturalism”, which is the philosophical viewpoint that begins with the realisation that we humans (our thought and talk included) are surely part of the natural world. The key claims of the chapter are that subject naturalism is importantly prior to object naturalism, because the latter depends on an assumption about language that might prove false, from the former’s perspective; and that there are good reasons for thinking that the threat is a serious one – that representationalism is might well turn out to be a bad (proto)-theory. If so, then the matching game turns out to be a bad analogy for the task that confronts a philosophical account of the place of language in the natural world.

But what alternative is there? Many of the remaining essays in this collection offer sketches of an answer to this question. In most of them, my role feels to me something like that of a real estate agent, making brief visits to a neglected property, with various kinds of reluctant clients in tow. Usually, the imagined clients are proponents of some rival theory, and I am attempting to convince them that their needs would be better met, at lower cost, by moving to this alternative.

Most of these rival theories are associated with one or other of various familiar approaches to the puzzle of the matching game – to the problem that we seem to have a lot more true statements than naturalistically respectable truthmakers. The best way to get a sense of where my alternative fits in is to begin there, with a brief survey of the more familiar options.

1.2 Placement strategies

The problem is that of “placing” various kinds of truths in a natural world.¹ We seem to have more truths than truthmakers – more stickers than places to put them. Since that puzzle thus turns on an apparent mismatch between the cardinality of two different sets, it should come as no surprise that there are three basic kinds of solution. One argues that the two sets can be matched, just as they are; that there is some non-obvious mapping that does the trick. The second argues that the problem arises because we have undercounted on the right, and that there are actually more truthmakers available than we thought. And the third argues that we have overcounted on the left, and that actually there are fewer statements in need of truthmakers than we thought.

The first option can be called *reductionism*. A noteworthy recent version of this approach is an account due to Frank Jackson, now commonly called the Canberra Plan.² I contrast my approach to Jackson’s in two essays in this collection, ‘Naturalism and the Fate of the M-Worlds’ (Ch. 6), and ‘The Semantic Foundations of Metaphysics’ (Ch. 12). In the latter paper, especially, I try to exhibit the way in which Jackson’s program depends on substantial assumptions about language – in effect, Representationalism – and I argue that these assumptions are problematic, in various ways.

The second option is often felt to embrace two sub-options. One accepts the constraint imposed by Naturalism, but argues that there are more facts within the scope of natural science than we thought.³ The second argues that the constraint itself is at fault, and that we need to recognise that there are non-natural facts.⁴ It is debatable whether the distinction between these two sub-options is more than merely terminological – an issue as to what we call science – but we need not discuss that here. What is relevant is what the sub-options have in common, viz., that they attribute the original puzzle to excessive parsimony in our initial assessment of the available truthmakers on

¹I’m not sure where the terminology originates. It is used in this way by Blackburn (1993b), writing about an earlier discussion by McDowell.

²This label was first coined by John Hawthorne and me, in a predecessor of our joint paper ‘How to Stand Up for Non-Cognitivists’ (Ch. 5), originally presented at the Auckland meeting of the Australasian Association of Philosophy in 1994. As we explain in the published version, the original intent was ironic, though the label became a badge of honour for our Canberra opponents.

³David Chalmers’ (1996) view of consciousness is a familiar example.

⁴The classic example is Moore’s (1903) view of moral facts.

the right hand side of the model. In the earliest essay in this collection (Ch. 2, ‘Metaphysical Pluralism’), I contrast these “super-additive” or “expansionist” approaches (as I call them there) to what I argue to be a more economical way of doing justice to similarly pluralistic intuitions. The alternative, which in the essay I associate with Wittgenstein, attributes the plurality in question to diversity in the things we do with language, rather than to some pre-existing metaphysical or natural diversity in the world. Most of the essays in this volume can be seen as engaged in the elaboration and defence of this “Wittgensteinian” proposal, in various ways; although only one piece (Ch. 10, ‘Immodesty Without Mirrors – Making Sense of Wittgenstein’s Linguistic Pluralism’) specifically addresses its Wittgensteinian credentials.

The third option, not surprisingly, is to try to reduce the size of the set on the left – in other words, to try to reduce the number of statements we take to require truthmakers. In this case, there are several sub-options, and it will be worth our while to distinguish them with some care.

Eliminativism. Recall that the stickers on the left of the model are supposed to represent true statements. An eliminativist deals with the excess – i.e., with the embarrassing residue, after all the obvious candidates are assigned to their naturalistically-respectable places on the right – by saying that we are victims of large scale error. Large subclasses of the statements we take to be true are actually systematically false. For example, perhaps there simply are no moral facts. If so, then all claims whose truth would depend on the existence of such facts are systematically in error.

Fictionalism. A similar but slightly less drastic view offers the same diagnosis of the apparent mismatch between statements and truthmakers, but with a more irenic conclusion. Eliminativists are inclined to compare the false statements in question to the claims of discarded scientific theories, and to recommend that they be accorded a similar fate. Fictionalists are more mellow about falsehood. They embrace the idea of “useful fictions” – language games in which false claims serve some useful purpose. The practices of making moral or modal claims might be beneficial in some way, for example, despite that fact that the claims concerned are not literally true. If so, we do not need to find truthmakers, but nor do we need to dispense with language games in question.

Expressivism. The same lesson – viz., that the “point” of some of the statements on our initial list is not to match worldly facts – is carried a stage further by expressivists.⁵ Expressivists maintain that some of the uses of language that we take to be statements are not genuine statements at all, but rather utterances with some other point or function. The suggestion is that once these pseudo-statements are pruned away, the apparent

⁵This is not to suggest that expressivism is a descendant of fictionalism. It might be more accurate to say that fictionalists are proto-expressivists, who have not yet realised that there is a live alternative to Representationalism.

imbalance between true statements and worldly truthmakers is eliminated, or at least reduced. The usual version of the puzzle rests at least in part on a kind of mistake about language, in the expressivist's view.

At this point, it is worth noting an important difference between fictionalism and expressivism. To make things concrete, consider the moral case. A fictionalist thinks that moral claims have an everyday use and a literal use. Taken literally (and interpreted as a moral claim), the statement "Harming children is worse than harming dogs" is false. (Why? Because, literally speaking, there are no moral facts to make it true.) Taken in its everyday sense, however – within the fiction in which we all participate – it may well be correctly said to be true.

In contrast, an expressivist has no need to admit that there is any sense in which such a statement is literally false. On the contrary, says the expressivist, taking it to be literally false is making a mistake about what kind of speech act it is. It is not the kind of speech act that *has* a literal truth-value, in the sense that the fictionalist intends.

As a result, an expressivist might hope to agree with everyday moral claims, without having to take anything back – without having to admit (even if only in her study, as it were) that all such claims are literally false. She agrees full voice with the everyday folk, and argues that the attempt to raise further issues – Are there *really* any such facts? – rests on a mistake about language. Once we see that moral claims are not genuinely descriptive, we see that such metaphysical issues rest on a category mistake. See things properly, the expressivist assures us, and you see that they simply don't arise.

1.3 Quasi-realism and globalisation

It might seem that the advantage of not having to say that our moral claims are literally false comes with a countervailing disadvantage. Does the expressivist not have to give up on the idea that there could be some everyday sense in which such a claim is true? Indeed, how is the expressivist going to account for the fact that we call such claims true and false, if they are not really in the business of making claims about how things are?

These issues are best addressed by the version of expressivism called *quasi-realism*, championed over many years by Simon Blackburn.⁶ Quasi-realism begins where expressivism begins, with the thought that the primary function of certain of our (apparent) statements is not that of describing how things are. But it aims to show, nevertheless, how such expressions earn a right to the trappings of descriptive "statementhood" – in particular, the right to be treated as capable of being true and false.

Blackburn emphasises that the appeal of quasi-realism is that it provides a way of dealing with some of the hard placement problems – the case of moral and aesthetic discourse, for example – without resorting either to implausible metaphysics or the

⁶See especially Blackburn (1984) and the papers collected in Blackburn (1993a).

error theory. If successful, quasi-realism explains why the folk practice of making moral claims is in order just as it is, and explains why further any metaphysical enquiry about whether there are *really* moral facts is inevitably missing the point (in being premised on a mistaken view of what we are doing with moral language).

Quasi-realism is important, in the present context, because the view proposed here can be thought of, in most respects, as a generalised or “global” version of quasi-realism. This way of locating the view is mentioned in many of the essays, and most explicitly in one of the most recent (Ch. 11, ‘Pragmatism, Quasi-realism and the Global Challenge’). To understand how the generalisation proceeds, note first that what expressivism does is to remove some (apparent) commitments from the matching game – to say that the matching model is a bad model of the relation of those commitments to the world. (What quasi-realism in particular adds is an account of why, on the surface, it “looks as if” the matching model is applicable.) In place of the matching model, presumably, expressivism offers some positive account of the use of the parts of language in question – some account compatible with the basic (“subject naturalist”) premise that the creatures employing the language in question are simply natural creatures, in a natural environment.

Typically, of course, expressivists do all of this *locally*. They think that some of our claims are genuinely factual, or descriptive (and hence, presumably, characterisable in terms of the matching model, in so far as it works at all). And they think that for any of our claims or commitments, there is a genuine issue whether it is really factual, or descriptive. (The expressivist’s alternative is needed when the answer is “No”.) However – and this is a crucial point – the belief that there is such an issue, and the belief that some claims are genuinely descriptive, play *no role at all* in the positive story, in the case of the commitments the expressivist regards as not genuinely descriptive. In other words, the expressivist’s positive alternative to the matching model does not depend on the claim that the matching model is *ever* a useful model of the relation between natural language and the natural world. So there is no barrier, in principle, to abandoning the matching model altogether, and becoming a *global* expressivist.

In essence, this global expressivism is the view that I want to defend. I want to defend it in a version that takes over from quasi-realism a strong emphasis on questions about why parts of language that begin life as expressions and projections should take the form that they do – why, in particular, they should be declarative in form, and capable of being regarded as true or false. But as I will explain, I part company with the quasi-realist on some of the details – necessarily so, perhaps, if the quasi-realist’s answers depend on the idea of emulating “genuinely representational” claims, because for me there are no such things.⁷

⁷As I will explain, I think that quasi-realism’s commitment to the idea that some statements involve genuine representation has hampered the enterprise of developing an adequate general theory of judgement

1.4 Naturalism without representationalism

Approached from this direction, then, the view I want to defend can be regarded as a kind of global expressivism. But it differs from local varieties of expressivism not simply in doing globally what they do only locally, but in a more fundamental respect. Local versions of expressivism accept Representationalism in some domains. Their message is simply that the matching game is not as widely applicable as we tend to assume – some of our statements (or apparent statements) have other, non-representational functions, and hence are not in need of truthmakers. We might imagine an expressivist of this sort raising the possibility of a language in which, as a matter of fact, all the (apparent) statements had this kind of non-representational function. This would be to imagine a language (perhaps even our own language) for which, contingently, a globalised expressivism does turn out to be the right story. But it would be to imagine it while keeping in play the proto-theory, and the notion of genuinely representational language.

I want to go a stage further. I am not proposing merely that genuine representation turns out to be a linguistic function that is not in play in our own language, but that representation (*in this sense* – more on the importance of this qualification in a moment) is a theoretical category we should dispense with altogether. The right thing to do, as theorists, is not to say that it turns out that none of our statements are genuine representations; it is to stop talking about representation altogether, to abandon the project of theorising about word–world relations in these terms. It is a bit like the familiar case of simultaneity: the lesson of relativity is not merely that we live in a world in which absolute simultaneity does not make itself manifest, but that we should abandon the notion of absolute simultaneity altogether, for theoretical purposes.

If representation goes by the board in this way, then there is a sense in which expressivism triumphs by default. After all, the defining characteristic of expressivism is that it offers a non-representational account of the functions of some part of language – and if representation goes by the board, there is nothing else left. Still, there is another sense in which the term expressivism is a little unhappy, in this context, because of its associations with familiar local forms of expressivism, which take for granted a representationalist framework. Sometimes, therefore, I call the view “pragmatism” – though this term, too, has mixed associations.

and assertion. If nothing else, it has obscured the explanandum, by disassociating the issue as to why moral claims (say) take the form that they do from the deeper question as to why any claims take such a form. Far better, in my view, to begin further back, tackling the deeper question in a manner that does not simply presuppose Representationalism – here Brandom offers an explicit model of the methodology required, I think (see Ch. 10, §8 and Ch. 14) – and then hoping to sweep up the problem areas as special cases. ‘Semantic Minimalism and the Frege Point’ (Ch. 3) offers a similar methodology for responding to the Frege-Geach argument. In each case, the quasi-realist’s tactical error is to be too charitable to his opponent, in allowing that the relevant account of genuinely representational statements is already in hand. (More on this in §1.8 below.)

For the moment, I want to emphasise two things about this view, whatever it is called. First, there is a clear sense in which it is naturalistic: it adopts the scientific perspective of a linguistic anthropologist, studying human language as a phenomenon in the natural world. It may reject Naturalism, or “object naturalism”, but its own naturalistic credentials are not in doubt.

Secondly, the view does not claim, absurdly, that there is nothing to be said about the relation of our words to the natural world. On the contrary – and as the example of local forms of expressivism makes abundantly clear – it is likely to allow that there is much to be said about why natural creatures in our circumstances come to use the forms of language in question. What it denies is simply that *representation* turns out to be a useful theoretical category, for saying what needs to be said about word–(natural-)world relations. (This is compatible not only with there being other theoretical vocabularies for characterising word–world relations, but also – as we shall see in a moment – with the notion of representation having other, more useful, theoretical applications.)

In one sense, this is a familiar idea. There are famous critics of Representationalism in modern philosophy, such as Dewey, Wittgenstein and Rorty. In another sense, it is a view which can be very hard to bring into focus. For my part, I have long felt that it occupies a peculiar location on the contemporary philosophical map: in one sense almost central, or at least easily accessible from familiar and popular places; in another sense almost invisible, almost unvisited. Part of the reason, presumably, is the strong intuitive appeal of the simple model of language for which the matching game is a metaphor. Representationalism can easily look obvious – more on this in §1.12 below – and is deeply embedded in contemporary philosophical theory.

Putting the position on the map, and revealing its virtues and accessibility, is a matter of visiting familiar locations that actually lie close by, and then calling attention to the paths that lead in the right direction. (I have already noted the path from quasi-realism.) In a sense, however, we need to do this from several directions simultaneously. Each path individually can easily seem to be obstructed from another angle. Let us look briefly at some of the angles, and at how they are tackled in the essays in this volume.

1.5 Avoiding the wrong kind of pragmatism

Pragmatists and expressivists point out that various kinds of commitments seem to have distinctive links to aspects of our own psychology, or to contingent features of our situation more generally. For example, evaluative, probabilistic and causal commitments have all been held to be distinctively linked to (various aspects of) the fact that we are decision-makers and agents. Facts of this kind – facts about the kinds of creatures we are, and about the relevance of these characteristics, in our relation to our natural environment – are an important part of the raw material for the account that my kind of

pragmatist wants to give of the functions and genealogy of particular parts of language. One of the reasons why the representationalist model is a bad theory, the pragmatist wants to say, is that it does not pay enough attention to these factors. It is blind to the “located” character of various bits of language – to their dependence on various contingent features of the circumstances of the natural creatures who use them.

One factor hiding my kind of pragmatism from view, at this point, is that there are views which claim similar sensitivity to the contingent dependencies of language, without ever leaving the familiar comforts of Representationalism. In recent literature, many views of this kind avail themselves of the notion of “response-dependence” – a notion claimed by at least one of its chief proponents (Johnston 1993) to be a step in the direction of pragmatism. Not much hangs on the label, but it is important to see that there is a very different way of giving theoretical voice to similar intuitions about the relation of language to contingent aspects of speakers’ circumstances. ‘Two Paths to Pragmatism’ (Ch. 3) draws this crucial distinction, and argues the cases for the non-representational path.

In my view, the recognition of this kind of contingency-to-speakers’-circumstances, on the one hand, and of the possibility of non-representationalism, on the other, fit together very naturally, and indeed reinforce one another. However, we really need to focus our attention in two places simultaneously, to see the benefits – if we look at one issue or other independently, ignoring the possibility of an unconventional approach to the other, we miss the attractions in question, being blind to the mutual benefits.

1.6 Keeping the lid on metaphysics

Like expressivism in general, quasi-realism motivates a kind of metaphysical quietism about the domains to which it is applied. Given that the commitments in question are not genuine factual commitments, metaphysical questions about (say) whether there are *really* any moral facts are simply misguided. The quasi-realist maintains that they involve a kind of category mistake, a misuse of moral language. However, since this case for quietism rests on the distinction between genuinely factual and not genuinely factual commitments, one might worry that it is unavailable, when that distinction goes by the board.

This concern is actually unwarranted, in my view. The reason the distinction goes by the board, in my global kind of expressivism, is that Representationalism itself goes by the board. So no commitments at all are treated as genuinely factual, in that sense, because the theoretical category is no longer on the table. But the expressivist’s argument for metaphysical quietism depended only on the fact that it was not on the table, in particular cases; and hence simply generalises, if it is never on the table.

Still, the more friends the better, and in several of these essays I appeal to the au-

thority of some famous allies, in support of the kind of metaphysical quietism that my view requires and entails. Carnap is one of these allies, but his celebrated attack on metaphysics, in ‘Empiricism, Semantics and Ontology’ (Carnap 1950), is often thought to have been decisively rebutted by Quine (who went on to make the world safe again for metaphysics, according to a popular version of the history of twentieth century philosophy). This reading is quite misguided, in my view. Not only do Quine’s criticisms not touch Carnap’s metaphysical quietism, or indeed his pragmatism; but Quine, too, should be read as a quietist and a pragmatist – in most respects, indeed, a more thoroughgoing one than Carnap himself. Ch. 13 deals directly with these matters.⁸

My debt to Carnap is not entirely one-sided, however. The general view defended in these papers offers significant new support to Carnap’s position, against Quine’s objections. Quine challenges Carnap’s entitlement to put fences around linguistic frameworks – to maintain, in effect, that there are several existential quantifiers, each doing different duty in a different framework. Quine argues that once we abandon the analytic–synthetic distinction, the fences disappear. We are left with a single arena, as it were, and a single existential quantifier, bullishly surveying the whole. I think that the key to resisting this objection is to make explicit something which is only a kind of implicit corollary in Carnap’s own work, viz., the idea that what distinguishes linguistic frameworks are the kinds of functional and genealogical factors to which expressivists call our attention. Carnap ought to say that the pluralism of linguistic frameworks is a functional pluralism, in this sense.⁹

As I note, this kind of functional pluralism challenges a kind of mono-functional conception of language that seem implicit in Quine’s own view – for Quine, *the* significant task of the statement-making part of language is that of recording the conclusions of an activity which is ultimately continuous with natural science. One of the interesting things about the role of this kind of methodological monism is how uncomfortably it sits with other aspects of Quine’s views about language. In particular, as other writers have noted, it is hard to reconcile with his deflationary views about truth and reference. In this respect, too, as we are about to see, Quine really ought to count as an ally of the kind of program I have in mind – an ally who provides his own answer, in effect, to his famous challenge to Carnap’s pluralism.

⁸Concerning the interpretation of Quine, see also Ch. 1 and Price (2007a).

⁹As I explain in Ch. 13, however, there is an important sense in which Carnap does not need to deny that there is a single existential quantifier – far better to say that we have a single logical device, with a variety of functionally distinct applications. I note that this formulation also defuses Quine’s well-known objection to Ryle’s pluralism about existential quantification. It also provides a model for a more general feature of the kind of view I want to recommend, viz., that it combines unity or homogeneity at one level in language with diversity at another level. The higher-level forms and structures thus become multi-purpose devices, in a novel sense. More on this in §1.9 below.

1.7 Keeping the lid on semantics

As the matching game itself illustrates, quietism about metaphysics needs to go hand in hand with quietism about semantics. The game is a metaphor for a linguistically-grounded methodology that has come to dominate contemporary metaphysics. Doing metaphysics this way, one begins with statements we take to be true, and then ask what *makes them true*, or to what their terms *refer*.

However, it is important to notice that there are two ways to take such enquiries, a weak way and a strong way. In the weak way, the semantic terms involved can be understood in a deflationary manner, and what is involved is merely semantic ascent, in Quine's sense. It "looks as if" we're talking about language – asking serious theoretical questions about the semantic relata of sentences and terms – but really we're just talking about the objects. Asking "What makes it true that snow is white?", or "What makes 'Snow is white' true?", is just another way of asking what makes snow white – a reasonable question, in this case, but a question to be answered in terms of the physics of ice and light, not in terms of the metaphysics of facts and states of affairs. There is no additional *semantic* explanandum, and no distinctively metaphysical question.

In general, then, this weak, deflationary view of the semantic terms allows us to read "What makes is true that P?" as something like "Why is it the case that P?", or simple "Why P?" This is a simple, first-order request for an explanation, which makes various and varied kinds of sense, depending on the subject-matter concerned. It is not a second-order enquiry revealing a theoretical commitment to a univocal or substantial relation of "truthmaking", of the kind required for linguistically-grounded metaphysics. Our theoretical gaze never leaves the world.

Moreover, the term "world" here need not be read as "material world". We saw how the combination of Representationalism and the (reasonable) assumption that language is a natural phenomenon leads Naturalists to want to find natural truthmakers for all true statements. But with Representationalism and that notion of truthmaking out of the picture, here – with all our semantic notions suitably deflated – we can ask "What makes is true that P?" with our gaze on other kinds of matters. We can ask "What makes it true that causing unnecessary harm to animals is wrong?", for example, requesting some sort of moral explanation or elucidation, without feeling any of the Naturalists' pressure to read this as an enquiry about the material world (or, for that matter, *metaphysical* pressure about some other kind of world).¹⁰

¹⁰Simon Blackburn makes a similar point about semantic ascent, construed in terms of Ramsey's redundancy theory of truth. Blackburn notes that on Ramsey's view, the move from 'P' to 'It is true that P' – "Ramsey's ladder", as he calls it – doesn't take us to a new theoretical level. He remarks (1998: 78) that there are "philosophies that take advantage of the horizontal nature of Ramsey's ladder to climb it, and then announce a better view from the top." In the present terms, the philosophies that Blackburn has in mind are those that fail to see that the fashionable linguistic methods – talk of truthmakers, truthconditions,

There is another use of weak or deflationary semantic notions which needs to be mentioned at this point, viz., their use in contexts of *interpretation*. If I say that “Schnee is weiß” is true iff snow is white, it might be said that there is a clear sense in which I am talking about language (in this case, German), and not about snow. For the moment, the main point to be made about these contexts is that they, too, don’t depend on substantial word–world semantic relations. As a result, once again, they are blind to the considerations that animate our Naturalists. The interpretative stance is blind to the kinds of distinctions marked by expressivists, for example. So even though it is true, in some sense, that in interpretation our gaze falls on language, it does not fall on the relation between words and the world, in the way that supports metaphysics – the way that depends on substantial semantic relations. (More on this in §1.11.)

So much for the case in which the use of semantic terms in metaphysics is read in a weak, or deflationary, sense. When the semantic notions are taken in the strong sense, however, then the metaphysician’s theoretical gaze can indeed rest on language. The resulting metaphysical program operates under the assumption that certain linguistic items – sentences, or terms, for example – have substantial semantic properties, or stand in substantial semantic relations. Given this assumption, questions posed about these semantic properties and relations can provide an indirect method of discovering things about the non-linguistic world.

In order to keep the lid on metaphysics – indeed, in order, more basically, to keep a lid on Representationalism itself, which combines with the Naturalist’s mantra to give rise to this kind of semantically-grounded metaphysics – it is necessary to keep the lid on the kind of substantial, non-deflationary semantics. Semantic deflationism, or minimalism, thus figures very prominently in the essays in this volume.

Here, too, of course, I rest heavily on the shoulders and authority of giants – in this case, the “giant minimalists” of the twentieth century, such as Ramsey, Wittgenstein, and Quine himself. But here, too, the debt is not entirely one-sided. The project as a whole offers a new view of the significance of semantic deflationism. It reveals some little-recognised advantages of theorising about language without substantial word–world semantic relations – in particular, that it leaves room for an attractive kind of pluralism about the roles and functions of linguistic commitments (precisely the kind of pluralism whose denial Naturalism presupposes, in effect). With this picture in view, the contrast between the apparent homogeneity of our talk of truth and the diversity of underlying function cannot help but make it implausible that semantic relations can be characterised in causal–functional terms, in any plausible way. As I say, I take this kind

denotations, and the like – add precisely nothing to the repertoire of metaphysics, unless the semantic notions in question are more robust than those of Ramsey, Wittgenstein and Quine. I am in complete agreement with Blackburn on this point, but I want to encourage him to walk his own plank: I think that the laying the semantic ladder horizontal defeats the vestige of Representationalism that still distinguishes his quasi-realism from my global view.

of consideration to provide significant support for semantic deflationism – support we miss if we miss the pluralism, and think of assertoric language as all doing the same kind of job.

1.8 The true role of truth and judgement

There are some respects in which I part company with familiar versions of semantic deflationism. In particular, I am sympathetic to the charge that the familiar disquotational versions of deflationism pay insufficient attention to the normative character of the notions of truth and falsity. Disquotational truth seems too “thin” to play its proper role in an adequate theory of the general features of assertion, commitment and judgement.

One motivation for this view, represented in this collection by my disagreement with Richard Rorty in ‘Truth as Convenient Friction’ (Ch. 8), is the feeling that pragmatists have often ignored the resources of their own theoretical standpoint – even, in a sense, their own principles – in seeking to equate truth with something like warranted assertibility. A better alternative, in my view, is to seek to explain *in pragmatic terms* why our notion of truth does not line up neatly with warranted assertibility – in other words, to explain what practical use we have for a stronger notion.¹¹ In my view, there’s a plausible answer to this question to be had in terms of the appropriate norms of assertion and commitment – in effect, roughly, the norms required to make an assertion be a *commitment*, rather than a mere expression of opinion.

This view puts the emphasis squarely on the normative character of truth, and like writers such as Wright (1992), I feel that a merely disquotational account of truth cannot do justice to this normative dimension. Unlike Wright, however, I do not think that this is a reason to abandon deflationism. On the contrary, as I have said, I think that what we need is a pragmatic, explanatory account of the role and genealogy of the distinctive linguistic norm in question – an expressivist account of a normative notion of truth.¹²

I had come to same conclusion from a different direction in earlier work. There, one motivation was an issue raised by Bernard Williams (1973): if truth is “thin”, or deflationary, why should its application be restricted to assertions. As Williams notes, it is not clear why such a notion of truth should not be used for endorsing other kinds

¹¹As the expressivist analogy ought to make clear, this needn’t amount to realism *about* truth, in some metaphysical sense; even if the notion of truth involved is strong enough to count as a “realist” notion, by the lights of someone who divides realists from anti-realists in terms of the kind of truth predicate each takes to be appropriate.

¹²In a paper not included in this collection (Price 1998), I discuss Wright’s dispute with Paul Horwich about these matters. I argue that although Wright is right about the need for normativity, deflationism wins the wider battle: as elsewhere, what we need is a deflationary, pragmatic or expressivist account of the functions and genealogy of the relevant norms. (The positive proposals made in that paper are developed at greater length in Ch. 8.)

of utterances, such as questions or requests.¹³ In *Facts and the Function of Truth* (Price 1988) I proposed an answer in terms of normative structures associated with assertion and commitment, and in particular their role in highlighting disagreement. I suggested that this makes much more sense for some of the things we do with language than for others, and that this is the major constraint on the bounds of assertoric language.

This proposal relied on a thesis about the pragmatic significance of norms of truth and falsity (or, rather, of conversational norms whose roles seem closely approximated, for us, by those of truth and falsity): roughly, the suggestion was that norms of this kind play an indispensable role in making disagreements “matter” to speakers – in helping to ensure that speakers who disagree do not simply “talk past one another” (like customers ordering different meals in a restaurant, to use one of my examples).

In *Facts and the Function of Truth* I argued that this view of truth throws new light on the kinds of intuitions that have often motivated expressivism and non-cognitivism, to the effect that some areas of discourse are less “factual” than others. The book begins with a sceptical examination of the fundamental distinction on which such views rely, between “cognitive” and “non-cognitive” uses of indicative utterances. There are many terms in use to mark this (claimed) distinction – “descriptive” versus “non-descriptive”, “belief-expressing” versus “non-belief-expressing”, “factual” versus “non-factual”, and so on. However, it seemed to me that they were all in the business of taking in each other’s washing, and that there was no well-founded distinction to be found.

This may sound as if it was an assault on the foundations of non-cognitivism, but my intentions were more even-handed. I took it that I was criticising a presupposition that non-cognitivists normally share with their opponents. Both sides presuppose that there is a genuine distinction in language, and disagree only about where it lies. I was arguing that this presupposition is mistaken, and hence disagreeing with both sides.

More importantly, my broader sympathies lay in the non-cognitivist camp. In effect, I was proposing what I am now calling global expressivism. And in the second part of the book, I offered a way of making sense of some of the guiding intuitions of non-cognitivism, in terms of the thesis just described about the dialectical role of truth. I suggested that one could understand such intuitions in terms of various ways in which disagreements might reasonably turn out to have a “no-fault” character (linked to a pragmatic or expressivist understanding of the functional role of the discourse in question). I argued that resulting classification is a matter of degree – the picture offers no sharp distinction between factual and non-factual uses of language.

In effect, then, a large part of *Facts and the Function of Truth* was an assault on Representationalism, along the vulnerable flank on which – even by its own lights – it

¹³Some writers are tempted to appeal to syntax at this point, arguing that it is “ungrammatical” to attach a truth predicate to anything other than an indicative sentence; but this surely prompts the same question about the grammar.

needs to mark its borders with non-representational uses of language. The assault was a kind of pincer movement. From one side, I tried to undermine the presupposition that there was a well-founded distinction to be found in the territory in question. From the other, I tried to account for some of the linguistic phenomena that might be considered relevant to the issue – especially the application of the notions of truth and falsity – in a way which led naturally to the conclusion that the distinctions in question were a matter of degree.

As I have explained, many of the papers in the present volume can be seen as making a more direct assault on the same target. This new attack scales the siege ladder provided by quasi-realism, exploiting semantic and metaphysical deflationism to argue that there is no inner citadel, no genuinely representational core that such a ladder cannot reach – nowhere we need anything but quasi-realism, in effect.

It is easy to be mistaken about where this route leads. As I note in several of the present papers (e.g., Ch. 9, §6), many writers have thought that semantic deflationism provides an easy victory for cognitivism, by making it a trivial matter that moral claims (say) are “truth-apt”. They fail to see that putting moral claims on a par with scientific claims need not be a victory for a Representationalist view of the former. It can be – and in this case, *is* – a defeat for a Representationalist view of the latter; a global reason for rejecting Representationalism itself. This mistake has contributed greatly to the near-invisibility of the global expressivist position in contemporary philosophy, in my view – though expressivists themselves must share some of the blame, in failing to see that the stable response lies at the opposite extreme, in a global version of their own position.

However, it seems to me that once the assault route is clearly in view – its true endpoint in sight – it must still be used with caution. For it rests a little too heavily on the horizontal character of Ramsey’s ladder, on the thin and universal character of semantic ascent, minimally conceived, and hence is in danger of obscuring some important and substantial matters about assertion and judgement. A conventional quasi-realist might hope that he leaves these matters untouched, at least within the representational core, where presumably they matter most; but that option is no longer available, if there is no such core. The more tortuous route of *Facts and the Function of Truth* had the advantage that it put one at least one aspect of these issues centre-stage, and proposed a way of reconstructing non-cognitivist insights within the resulting framework.

1.9 Two-layered language

As I note in *Facts and the Function of Truth*, my view of truth suggests a distinctive two-level picture of the functional architecture of truth-evaluable uses of language. At the higher level, the picture offers us a certain kind of unity, or univocity: truth is essentially the same conversational norm, in all its core applications. (Contrast this to

the *local* quasi-realist's distinction between genuine and "quasi" truth.) At the lower level, however, there is room for a multiplicity of functions – a multiplicity of linguistic tasks or "games", each associated with different aspects of our psychology, needs and situation.

This two-level picture has many attractions, in my view. By prising apart word–natural-world relations from what everything common to the higher level (i.e., from the resources needed for a single, unified account of assertion, commitment and judgement), it adds a new dimension to linguistic theory, a new degree of freedom for functional variability. The effect is to open up regions of theoretical space that are simply invisible, when these two kinds of factors are squashed together. (Here, think of a child's pop-up book. As we open the page, the model lifts into view, transforming a flat jumble into something with recognisable three-dimensional structure: here is the Opera House, and there is the Harbour Bridge.)

Once again, orthodox forms of expressivism, and quasi-realism in particular, can be seen in hindsight as attempts to occupy these regions, working their way laboriously around the limitations imposed by the lack of the missing dimension. When the two levels are prised apart, and the model expands into the new dimension thus made available, these attempts fit naturally into place, without all the distortion.

The model offered in these essays is admittedly sketchy, at both levels. But at both levels, much of the necessary work has already been done elsewhere. The components are available off the shelf, as it were, as ready-made products of familiar projects in contemporary philosophy. Indeed, what is mainly novel about my proposal is simply the idea of connecting these projects together in this way.

At the lower level, the project in question is that of Blackburn, and other contemporary expressivists in the Humean tradition. As I have said, I think that many of the insights of that tradition plug straight into my framework, once re-oriented in the way I recommend.

At the higher level, my proposal about the global conversational role of truth cries out for incorporation into a broader account of the nature and genealogy of assertion, commitment and judgement – an account built on foundations that do not presuppose Representationalism, of course. I sketch a proposal as to how such an account might go in a couple of the papers in this collection (e.g., Ch. 10, §9), as well as in *Facts and the Function of Truth*, but more certainly needs to be said. But here, too, I am optimistic about what is already available. I think that much of what I need is to be found in the kind of inferentialist account of assertion developed in detail by Robert Brandom (1994, 2000). The crucial thing, from my point of view, is that Brandom explicitly rejects a Representationalist starting point, offering, as he puts it, an "expressivist alternative" to the "representational paradigm". (Brandom, 2000: 10)

There may seem to be a tension in linking my proposal to these two different "expressivisms": Brandom's, on the one hand, and Humean expressivism, on the other.

After all, expressivism in the Humean sense relies on “world-tracking” conception of genuine, full-blooded assertion, offering its treatment of particular cases by deliberate contrast with assertions so construed. Whereas Brandom, as just noted, explicitly rejects such a view of assertion. My proposal resolves this tension by abandoning the Representationalist residue in Humean expressivism. With this gone, we are free to help ourselves to Brandom’s view, as an account of what I have called the higher level in a two-level picture of assertoric language – in other words, an account of what all assertoric vocabularies have in common. And this is entirely compatible with also adopting off the shelf, as an account of the lower level, much of what Humean expressivists have to say about the functions of particular concepts and vocabularies.

So there is no real tension here – quite the contrary. But the apparent tension is revealing, in my view. It relies on a bifurcation in the conceptual territory surrounding the notion of representation in contemporary philosophy, which is worth noting and making explicit. (Indeed, I think that much of the appeal of Representationalism rests on a failure to make it explicit.)

1.10 Two notions of representation

Consider, then, the notion of representation, type or token, as it is used in cognitive science, and in contemporary philosophy of language and philosophy of mind. (Imagine a survey of these fields.) My proposal is that we can usefully distinguish two nodes, or conceptual attractors, around which the various uses tend to cluster. One node puts the system–world link on the front foot. It gives priority to the idea that the job of a representation is to *covary* with something else – typically, some *external* factor, or environmental condition. The other node gives priority to the *internal* cognitive role of a representation. A token counts as a representation, in this sense, in virtue of its position, or role, in some sort of cognitive or inferential architecture – in virtue of its links, within a network, to other items of the same general kind.¹⁴

In the grip of Naturalism, one naturally assumes that these two notions of representation go together; that the prime function of representations in the internal sense is to do the job of representing in the external sense. It takes some effort to see that the two notions might float free of one other, but it is an effort worth making, in my view. The vista that opens up is the possibility that representation in the internal sense is a much richer, more flexible and more multipurpose tool than the naive view always assumes.¹⁵

¹⁴I develop this distinction at greater length in Price (2008), calling the two notions *i-representation* and *e-representation*, respectively.

¹⁵Once again, quasi-realism provides a useful stepping-stone. The quasi-realist is already committed to the idea that something can behave for all intents and purposes like a “genuine” belief, even though it has its origins at some “non-cognitive” level.

Once the distinction between these internal and external notions of representation is on the table, it is open to us to regard the two notions as having different utilities, for various theoretical purposes. In particular, it is open to us to take the view that at least by the time we get to language, there is no useful external notion, *of a semantic kind* – in other words, no useful, general, notion of relations that words and sentences bear to the external world, that we might usefully identify with truth and reference. This is the conclusion that a semantic deflationist has already come to, from the other direction, as it were. On this view, the impression that there are such external relations will be regarded as a kind of trick of language – a misunderstanding of the nature of the disquotational platitudes. But we can think this without rejecting the internal notion: without thinking that there is no interesting sense in which mental and linguistic representation are to be characterised and identified in terms of their roles in networks of various kinds.

Networks of what kinds? We may well want to distinguish several very different conceptions, at this point. According to one conception, the relevant kind of network is causal (or causal–functional) in nature. According to another, it is normative and inferential. According to a third, at least arguably distinct from the other two, it is computational.¹⁶ But however it goes, the notion of representation involved can be divorced from any external notion of representation, thought of as a word–(natural-)world relation of some kind.

The sticker metaphor is useful again at this point. Think of these internal notions of representation as offering an account of what gives a sticker its propositional shape; what makes it the particular sticker that it is. As just mentioned, there are various possible versions of this internal account, but let us focus on the causal version, for the moment,

¹⁶In this case (as I am grateful to Michael Slezak for pointing out to me) Chomsky provides an excellent example of someone who not only thinks of representations in this way, but is explicit that it need no be accompanied by a referential conception:

As for semantics, insofar as we understand language use, the argument for a reference-based semantics (apart from an internalist syntactic version) seems to me weak. It is possible that natural language has only syntax and pragmatics; it has a “semantics” only in the sense of “the study of how this instrument, whose formal structure and potentialities of expression are the subject of syntactic investigation, is actually put to use in a speech community”, to quote the earliest formulation in generative grammar 40 years ago, influenced by Wittgenstein, Austin and others [Chomsky, 1975, Preface; 1957, 102-3]. In this view, natural language consists of internalist computations and performance systems that access them along with much other information and belief, carrying out their instructions in particular ways to enable us to talk and communicate, among other things. There will be no provision for what Scott Soames calls “the central semantic fact about language, . . . that it is used to represent the world”, because it is not assumed that language is used to represent the world, in the intended sense (Soames 1989, cited by Smith 1992 as the core issue for philosophers of language). (Chomsky, 1995: 26–27)

for definiteness. And let us make explicit that the line drawing on the right-hand page depicts the world as seen by natural science.

The first possibility we need to call into view is that there may be a lot more stickers given shape by their internal causal roles than stickers whose truthmakers may be found on the right hand page. Once again, quasi-realism is helpful to keep in mind at this point. Presumably, a quasi-realist maintains that our “quasi” beliefs – about morality, chance, or whatever – play very much the same roles in our internal cognitive economy as genuine beliefs (being distinguished mainly by characteristic *additional* functional links, in this case to action). Thus they are well-shaped stickers, despite matching no corresponding shape in the natural world.

Once we have reached this stage, we can progress to the mature view. The key step is a shift in our conception of our theoretical goals, a shift from the project of matching stickers to shapes in the natural world to the project of explaining (in natural world terms), how stickers obtain their characteristic shapes. Freed of the requirement that they must bear semantic relations to the natural world, stickers – or representations in the internal sense – can now occupy a new dimension of their own in the model, orthogonal to the natural world. Like the figures in our pop-up book, they stand up from their bases in the natural world, without being constrained to match or resemble anything found there.

Of course, a pop-up book does all the work for us, as we open the page. For a more illuminating metaphor, let us make the construction into a puzzle: a sort of three-dimensional jigsaw puzzle. We begin with a large collection of shapes or pieces, each one a statement we take to be true, and a large board or playing surface, depicting the natural world (in such a way as to give prominence to our own situation, as creatures with certain attributes and situation, within that world). In effect, our task is then to solve two kinds of puzzle simultaneously. We need to arrange subsets of the pieces into clusters, fitting them together so that, as in an ordinary jigsaw puzzle, the shape of each is defined by its relations to its neighbours (and eventually, perhaps, to the super-cluster of all the pieces). And we need to position each of the resulting clusters in the correct place on the board as a whole, so that its edges bear the right relations to particular features of the situation of the speakers (ourselves, in this case) who are depicted on the board.

Consider the pieces representing probabilistic statements, for example. They need to bear certain internal relations to one another, corresponding to the inferential or causal–functional links that define internal representations and their conceptual components in general. But they also need to bear the appropriate functional relations to the decision behaviour depicted on the underlying board, in order to count as *probabilistic* statements at all. (In this case, in fact, the latter constraint is likely to involve an additional complexity. Roughly, the development of probabilistic concepts seems to enable, or perhaps go hand in hand with, a distinctive kind of cognitive architecture –

let it be the Bayesian model of belief revision, for the sake of the example. This means that fitting the probabilistic pieces properly into the jigsaw requires that they display the appropriate alignment not simply with the behaviour depicted on the board, but also with some general features of the architecture on view throughout the model as a whole.)

At least very roughly, then, the first stage of the puzzle is concerned with what makes a piece of the puzzle a *statement* at all. The second stage is concerned with pragmatic factors about its use that may play a crucial role in determining *what* statement it is – what its *content* is, as we would normally say. Missing altogether is the idea that the latter fact is determined by some matching to a shape already discernible in the natural world.

The upshot is a model in which there is a substantial *internal* notion of representation – a substantial theory as to what gives a piece or a pop-up figure its shape – but no substantial *external* notion of representation. As the model illustrates, moreover, internal notions of representation are not constrained by the cardinality of the natural world. So long as we find a role for pieces which is not that of matching outlines in the natural world, we can happily allow that there are many more pieces than available outlines. In effect, this is the original insight of expressivism and quasi-realism, here given a more attractive home, in a version of the picture in which external representation disappears altogether, for theoretical purposes.

Of course, the model still allows for a deflationary conception of the matching relation. Earlier, we thought of this conception by analogy to the version of the matching game in which the picture on the right was constructed by tracing around the outlines of the corresponding stickers on the left. Each statement identifies its own truthmaker, as Ramsey’s horizontal ladder requires, with natural truthmakers in no sense distinguished. This works just as well in the new, richer model, in which the pieces themselves are given their shapes as structured artifacts in a natural world. If we imagine taking a photograph of the whole construction, we will be able to match pieces to that image, one by one. But there is no special role for the natural world, in this case – indeed, there is a much diminished role for the natural world, compared to the project of explaining the shapes and structures of the model in the first place.

1.II Matching as interpretation

There is another “deflationary” version of the matching game, played with two sheets of stickers from different designers. Here the aim is to match the two sets, sticker by sticker. We try to place the Opera House sticker from the left on the Opera house sticker on the right – or at least on what we take to be the best candidate for an Opera House, in the strange and foreign graphical vocabulary of the stickers on the right.

This version of the game provides an analogy to the project of radical interpretation, as conceived by Quine and Davidson. I have introduced it to note that it, too, is blind to the cardinality difficulty, at least unless the two sheets concerned correspond to languages with very different conceptual resources. Matching sentences to sentences, the interpreter does not care about whether either matches some state of affairs in the natural world (unless it bears in some specific way on the available evidence, at least). Thus an interpretative notion of representation also counts as internal, in the relevant sense – a fact which goes hand in hand with the common view that the project of radical interpretation is compatible with deflationary views of truth and reference.¹⁷

The case of radical interpretation illustrates a way in which many of the concerns of philosophy of language seem relatively untouched by the kind of viewpoint I am recommending. Representationalism seems to play a less fundamental role than might be assumed at first sight. Consider a model-theoretic standpoint, for example. So long as we construct our models from the inside, as it were, formalising the structures visible to us already as language users (like a child drawing lines around the outlines of her stickers), the entire process can be blind to the cardinality concerns – blind to the considerations that drive the placement concerns. These concerns only arise if we attempt to transform our model theory into a view of the relation of language to the natural world, in the Naturalist's sense. That transformation rests on Representationalism, but the model theory itself does not.

In a similar way, much of the formal machinery of philosophy of language seems entirely compatible with the present viewpoint, so long as it is thought of in this internal way, as formalising and describing the structures and relations characteristic of language at the higher, homogeneous, level – the level which is blind to the underlying functional distinctions associated with the origins and roles of particular groups of concepts.

It may seem that a theory of this kind is bound to be a kind of sham. If I am right, after all, then it hides the underlying diversity, and does not speak to the interesting and various relations of language to the natural world. It is indeed bound to seem unsatisfactory, to some philosophical temperaments – to philosophers at home with the contemporary integration of metaphysics and philosophy of language; with the explicit attempt to study reality “through the lens of language”, via truthmakers, reference relations, and the like. To philosophers of this temperament I have offered the hypothesis that their viewpoint rests on a mistake about language and its place in the natural world – a deep, first-order, scientific mistake. (Again, this is just the familiar suggestion made by several generations of expressivists, but now cast in a more general and more stable form.) If the hypothesis is correct, then the approach of these philosophers (next to which the approach described above looks like a sham) is no real alternative at all.

My point is that reaching this conclusion need not mean throwing out the baby of

¹⁷Cf. Ch. 10, §5, and Williams (1999).

philosophy of language, along with the bathwater of metaphysics. Admittedly, it will not always be easy to say where baby stops and bathwater begins. Finding the boundary is a big project, to which the present conclusions are merely a preliminary. As the case of interpretation illustrates, however, we can be confident that the boundary is there to be drawn, somewhere short of the entire contents of the bathtub.

1.12 Is Representationalism obvious?

These distinctions help to clarify what is and is not entailed by a rejection of Representationalism. The guiding principle is that so long an apparently Representationalist intuition trades only in the deflationary semantic notions, we anti-Representationalists have no reason to reject it.

Clarity on this point can go some way to address the incredulity that tends to greet the denial that language is representational. Frank Jackson, for example, remarks that “[a]lthough it is obvious that much of language is representational, it is occasionally denied,” and goes on to observe that he has “attended conference papers attacking the representational view of language given by speakers who have in their pockets pieces of paper with writing on them that tell them where the conference dinner is and when the taxis leave for the airport.” Jackson asks how this could happen, and suggests “that it is through conflating the obviously correct view that much of language is representational with various controversial views.” (Jackson 1997: 270)

Jackson is quite right, in my view, that there is a danger of conflation, on one side of the case or other. For our part, we anti-Representationalists need to be clear that there is, indeed, a deflationary sense in which language conveys information. *Of course*, a sign affirming ‘P’ may inform us that P. But so long as deflationism is on the table, there is no more reason to think that this fact needs to be explained in terms of robust word–(natural-)world relations – the meat of big-R Representationalism, which is what we anti-Representationalists are opposed to – than there is in the case of the corresponding disquotational platitudes. And it is easy to muster intuitions to the contrary. Imagine a sign, in some bleak campus cafeteria, conveying to conference participants the following information: *Clients are forbidden to place inorganic waste in the green bins*. The sentence mentions a (quasi?) commercial relationship, a colour property, a prohibition, and some medium-sized dry(ish) goods: in one sense, then, at least four contributions to the information conveyed. Yet it can’t be *obvious*, surely, that the contributions are pragmatically univocal; that an explanation of the information conveyed by the sign will simply bottom-out at the level of content – at the level of differences between clients, colours, prohibitions and inorganic waste – leaving pragmatic factors no role to play?

Conventional Representationalism combines two assumptions about language and thought. The first (call it the “Content Assumption”) is that language is a medium for

encoding and passing around sentence-sized packets of factual information – the *contents* of beliefs and assertions. The second (call it the “Correspondence Assumption”) is that these packets of information are all “about” some aspect of the external world, in much the same way. For each sentence, and each associated packet of information, there is an appropriately shaped aspect of the way the world is, or could be – viz., the state of affairs, or fact, that needs to obtain for the sentence to be true.

Once both assumptions are in place, it is natural to regard language as a medium for representing these sentence-sized aspects of the external environment, and passing around the corresponding packets of information from head to head. My rival proposal rests on pulling the two assumptions apart, replacing the Correspondence Assumption with richer, more pluralistic, and non-semantic conception of the role of various kinds of linguistic information in our complex interaction with our environment. However, so long as we distinguish two notions of representation, as I recommended above, and allow the notion of information to go with the former, this proposal involves no challenge whatsoever to the view that language conveys information. About this, I agree with Jackson: it is obviously correct. My point is simply that it does not entail Representationalism, for it does not depend on the Correspondence Assumption – on the contrary, it is quite compatible with global expressivism.

1.13 A telescope for metaphysics?

In criticising Representationalism, I am arguing, in effect, that would-be metaphysicians need to play close attention to language, for two closely related reasons. First, as expressivists have long urged, issues that seem at first sight to call for a metaphysical treatment may be best addressed in another key altogether. And second, even if metaphysics is thought to provide the right key, it is doubtful whether anything is to be gained by arranging the score for semantic instruments. Concerning both points, one moral is that would-be metaphysicians cannot afford to ignore some deep issues in philosophy of language.

In one sense, then, I am very much on the same page with Timothy Williamson, who, in a recent paper about the role of linguistic issues in philosophy, criticises metaphysicians who believe they have nothing to learn from the philosophy of language:

Some contemporary metaphysicians appear to believe that they can safely ignore formal semantics and the philosophy of language because their interest is in a largely extra-mental reality. They resemble an astronomer who thinks that he can safely ignore the physics of telescopes because his interest is in the extra-terrestrial universe. In delicate matters, his attitude makes him all the more likely to project features of his telescope confusedly onto the stars beyond. (Williamson 2006: 127–128)

In another sense, however, I am on a very different page. For Williamson's central metaphor gives vivid and approving expression to the Representationalist conception, that I want to reject. As Brian Leiter glosses Williamson's point, in his introduction to the volume in which Williamson's paper appears, '[l]anguage is for the philosopher what the telescope is for the astronomer: the instrument by which the investigator makes contact with his "real" subject-matter.' (Leiter 2006: 6) With this conception in place, Williamson's and Leiter's lesson is a caution for philosophical practice. We should not take the capabilities of our telescopes for granted – 'instruments, unless they are well-understood, can corrupt our understanding of the "reality" to which they are our means of access.' (Leiter 2006: 6–7)

Presumably, the opponents Williamson and Leiter have in mind make one of two mistakes. Either they imagine that we can get at reality "directly", without the need of a linguistic lens; or they think that such a lens is so transparent that no issue of distortion can arise. In the present context, however, it should be clear that there is another opponent in the vicinity, to whom Williamson's metaphor simply turns a blind eye. For the metaphor compares language to a representational instrument; whereas, as we have seen, it is a recurring theme among linguistic philosophers that it is precisely the representational conception which leads philosophy astray, at least in certain cases.

This kind of opponent will endorse the sentiment that linguistic instruments, poorly understood, 'can corrupt our understanding of "reality"'. Indeed, in suggesting that Williamson's and Leiter's own view embodies a misguided conception of the functions of language, she will endorse the sentiment more vigorously than they themselves do. She will allow, presumably, that language is an instrument of some kind, an organ that we humans use in negotiating our physical environment. But she will insist that it does not follow that language is necessarily a *representational* instrument – at least in the manner, and cases, supposedly of relevance to metaphysics.

Of course, some versions of this position – familiar versions of expressivism, in particular – will differ at most in degree from the kind of view that Williamson has in mind. They will allow that language is at least sometimes "telescopic", even if the distortions and projections are more significant and wide-ranging than it seems at first sight. As I've emphasised, however, there's also a more systematic version of the view, holding that the representational conception is unhelpful *tout court*, at least as a tool for metaphysics. From this viewpoint, characterising language as the telescope of metaphysics looks like a symptom of a fundamental philosophical error.

But if language is not a telescope, then what is it? As Brandom points out, a traditional expressivist option is the lamp.¹⁸ I think that modern technology allows us to

¹⁸ Brandom notes that "to the Enlightenment picture of mind as a *mirror*, Romanticism opposed an image of the mind as a *lamp*. Broadly cognitive activity was to be seen not as a kind of passive reflection but as a kind of active revelation." (2000: 8). Brandom here cites M. H. Abrams' classic (1953) study of these two themes in the critical tradition. Abrams himself provides examples of the use of the lamp metaphor

make this a little more precise. Think of a data projector, projecting internal images onto an external screen. Even better, helping ourselves to one of tomorrow's metaphors, think of a holographic data projector, projecting three-dimensional images in thin air. This isn't projection *onto* an external, unembellished world.¹⁹ On the contrary, the entire image is free-standing, being simply the sum of all we take to be the case: a world of states of affairs, in all the ways that we take states of affairs to be.

At this point a newcomer (occupying a stance we ourselves cannot take up, perhaps) might notice that the projector contains an internal screen, the shapes on which match those in the external image, and conclude that the device is a telescope. Obviously, this gets things backwards, however. The facts seem to resemble the statements because the former are the projected image of the latter, not *vice versa* – the transparency is that of semantic descent, of Ramsey's ladder.

This may sound like a recipe for implausible idealism, so it is worth emphasising again its naturalistic credentials. The new model sits squarely within the project of understanding human linguistic usage, as a form of behaviour by natural creatures in a natural environment. But to see this, we need to be careful to distinguish the new model from the one we used to depict this explanatory project. The two models line up, more or less, on the left-hand side, where we find the page of stickers, or the projector's internal screen. Here, in each case, we have the raw data of linguistic usage, or at least some suitably selected subset of linguistic usage, such as the (apparently) assertoric uses, or the statements held true within the community in question.

But the two models differ on the right. In the new model, what lies on the right is the holographic image, a metaphor for the world in the most general sense, *as the language users in question take it to be* – the sum of all they take to be the case (and thus, as the metaphor has it, a projection of their usage). In the relevant naturalistic version of the original sticker model, however, we have something different on the right: the natural world, as viewed from the scientific standpoint – the context within which we

not only by romantic poets such as Coleridge and Wordsworth, but also by the "effusive" Edinburgh essayist, "Christopher North" (John Wilson). Wilson held the Chair of Moral Philosophy at Edinburgh, for which Hume had earlier been thought too irreligious, and seems to have shared Hume's antipathies, in this respect; *The Scotsman* called him a "mocker of the Scriptures", whose appointment would be "an outrage on public decency without parallel since Caligula made his horse a Consul." (Sharing Hume's tastes in another respect, he was a haggis-maker of some repute, whose recipe for "a sauce for meat and game" was preserved by Mrs Beeton for future generations.) Concerning the mind's lamp-like qualities, Wilson says that, as is "well known, . . . we create nine-tenths at least of what appears to exist externally . . . Millions of supposed matters-of-fact are the wildest fictions—of which we may mention merely two, the rising and the setting of the sun." (Wilson 1832: 721) For present purposes, I think we do better to build on Hume's rather more disciplined projectivism – even if updating his metaphor, too, in the way I shall explain, to cope with the demands of globalisation.

¹⁹This is how it differs from Hume's "gilding and staining" – a difference required by the global nature of the present view.

seek to *explain* the linguistic behaviour displayed on the left.²⁰

Crudely put – ignoring, for example, all the obvious grounds for holism – the explanatory project goes something like this. We find our speakers disposed to say “P” (i.e., ‘P’ appears in the list of statements on the left of the model). We now ask, “Why do they say that?”; and in general (without pretending that this distinction is sharp) we look for an explanation that refers both to features of the speakers, and to features of their natural environment. Note that in our own case, this attitude always looks sideways-on, or ironic. We say that P, and then wonder *why* we said so, how we came to be making a claim of that kind – looking for something deeper as an answer, of course, than merely, “Because we realised that P”. (This is the kind of irony characteristic of practitioners of the human sciences, of course, who cannot help but view themselves as examples their own objects of enquiry.)

In a somewhat more abstract form, the project becomes that of explaining the function and genealogy of our (internal) representations – that is, of the linguistic items we characterise as representations in whatever version of the internal sense (e.g., causal–functional, or inferential) we have in play. In the general case, as before, we expect explanatory contributions both from our speakers’ own nature and characteristics, and from their natural environment ; and a complex, relational story about the significance of the mix. (Again, think of Humean expressivism about value or causation as a model.) As I have emphasised, we should also expect to appeal to some general account of the role and significance of the cognitive or logical architecture, in virtue of which the items in question constitute representations, in the internal sense.

At least in the first place, then, this project regards language not as a telescope for metaphysics, but as the primary explanandum of a first-order scientific investigation: the project of understanding our own linguistic behaviour, and associated aspects of our cognitive lives, as characteristics of natural creatures in a natural environment. As we have known (in some sense) since the seventeenth century, this project requires that we consider our own contribution to our own commitments, as well as that of our environment. The project thus has a pay-off, closer to the traditional concerns of metaphysics. We get a less perspectival view of the world by subtraction, as it were, by asking “What kind of world looks like *this* from *here*?” – where the “this” refers to our commitments, to the holographic image, and the “here” to our natural situation, broadly construed.

In one sense, this sounds a lot like Williamson’s and Leiter’s project, but it has a rather more ironic flavour. We no longer see the philosopher’s task as to join the folk in peering *through* the linguistic telescope (contributing simply a heightened sensitivity to possible distortions). Rather, we treat its folk output as the starting point for a new enquiry, which has both anthropological and physical aspects. The theoretical

²⁰Note that this version of the naturalistic model is different in turn from the Naturalist’s version of the original matching game, where the task was to match stickers to *truthmakers* in the natural world.

framework is more rigorously naturalistic, in the explanatory sense, and by its own lights, much less naive. In particular, because governed by a more rigorous naturalism about ourselves, it is less willing to postulate metaphysical primitives, simply to allow us to take our own commitments at (what Representationalism takes to be their) face value. The default assumption is always that they should not be taken at their face value – that their face value is a sum of two components, one grounded in us, the other in the natural world.

1.14 The status of science

We might hope to be able to relax this ironic assumption, at some point. Eventually, perhaps, the human component might be entirely pruned away. We would have carved the world itself from the holographic image, as it were, and be left with a bare description of nature.

In practice, however, I think there are strong reasons for thinking that this limit is out of reach – that the contribution on our side never goes to zero. The deepest reasons are associated with the rule-following considerations. In my view, these considerations reveal that wherever generality matters, our judgements depend on dispositions to “go on in the same way” in some particular way. These dispositions might have been different, and hence constitute an ineliminable contingency in our linguistic standpoint. (These issues arise at several points in these essays, and most extensively in Ch. 7.)

What does this mean for science itself? Here, my proposal may seem to confront a dilemma. Either science is privileged, and non-perspectival, in which we seem to recover something like genuine representation after all (thus saving Representationalism, apparently, at least in this pared-down scientific form). Or we have to be ironic and perspectival about science, too; and that might seem to count as some kind of self-undermining, on science’s part.

Pushing towards the first horn, it might seem that natural science is automatically privileged in the kind of picture I have been describing. In the case of other frameworks, after all, the project is to explain in naturalistic terms why creatures of our kind, in our circumstances, have come to employ those frameworks: why creatures like us should have come to employ normative vocabulary, for example. In this case, the explanation is not going to appeal to pre-existing norms, presumably. Yet in the case of the frameworks we employ for talking about the natural world itself, things must surely be different – the natural entities themselves get cast in explanatory roles. Is the project not therefore committed to according a special status to these entities, by its own lights?

In my view, the right response to this challenge (see Ch. 6, §5) is to acknowledge that scientific ontology is privileged *from the standpoint of the present project* – How could it not be, since the project involves a first-order scientific reflection on our own linguistic

practice? – but to deny that this conclusion has the significance the objector claims for it. Science is only one of the games we play with language. Each game privileges its own ontology, no doubt, but the privilege is merely perspectival. Science is privileged “by its own lights”, but to mistake this for an absolute ontological priority would be to mistake science for metaphysics, or first philosophy.

Moreover, science certainly not privileged in the sense that the project cannot turn its spotlight on the language of science itself. On the contrary, I have already emphasised that the expressivism of my view is intended to be global, and in any case to cast light on such particular matters as the use of modal notions, of general concepts, and of a norm of truth. All of these matters seem at least as much in play in scientific language as in any other.

At this point, however, the second horn of the dilemma may seem threatening. If the language of science is not exempt, does the ironic viewpoint not become self-undermining, unable to take itself seriously? These are difficult issues, but I suspect that the threat is to a particular philosophical conception of science, rather than to the scientific enterprise itself. If we equate science with the perspective-free standpoint, the view from nowhere, then science so conceived is certainly under threat. But why not see this simply as a challenge from within science to a particular philosophical conception of science?

To make the case more concrete, suppose (as I have recommended elsewhere – see, e.g., Price 1991, 2001, 2007b; Menzies & Price, 1993) that we were to accept a perspectival genealogy for causation and related notions. This would be a scientific account of a particular aspect of human cognitive and linguistic practice, explaining its origins and function in terms of certain characteristics of ourselves (especially, the fact that we are *agents*, embedded in time in a particular way). A corollary would be that uses of causal concepts in science (including, indeed, in this very explanation) would be held to reflect the same agentive perspective. Some aspects of current scientific practice would thus be revealed *by science* to be practices that only “make sense” from this embedded viewpoint.

Would this be a fundamental challenge to science, a case of science undermining its own foundations? Not at all, in my view. On the contrary, it would simply be continuous with a long scientific tradition, in which science deflates the metaphysical pretensions of its practitioners, by revealing new ways in which they less god-like than they thought. Some of the greatest triumphs of science have been new ways of showing us how insignificant we are, from the world’s point of view; how idiosyncratic the standpoint from which we attempt to make sense of it. We humans may find these lessons unsettling, but science seems to have thrived on this diet of self-imposed humility. Why should the present case be any different?

True, this is hardly a decisive refutation of the kind of sceptical threat that may seem to lurk in the region of the second horn of our dilemma. It is fair to point out

that science has been thought to face such threats before, and pulled through in fine shape. But fair, also, to reply that this time it might be worse. The debate is admittedly inconclusive, but my kind of naturalist has one more card to play. Even if the recommended explanatory program were a deep threat to science, in the way that the sceptic imagines, that could hardly count as a reason for counting its conclusions *false* – deeply regrettable, perhaps, but that is a different matter altogether!

1.15 Naturalism without mirrors

The view proposed in these essays may seem radical, for it is certainly not popular, or well-marked in the contemporary philosophical landscape. In another sense, however, it involves no great departure from familiar ideas, for it is simply the place where some comparatively familiar paths turn out to converge. Indeed, my case for challenging Representationalism may be summarised in four simple steps. First, semantic deflationism already challenges Representationalism, by refusing to allow any substantial theoretical role or content to the semantic properties and relations. Second, the placement problems also turn out to challenge Representationalism, in the sense that expressivism and quasi-realism provide an attractive solution to some of these problems, but a solution of which a challenge to Representationalism turns out to be a consequence, when the considerations in question are pushed to their natural limits.

Third, arguments to the contrary at this point are either misdirected (I am thinking here of the attempt to use semantic deflationism against expressivism, which, as noted in §1.8, actually counts against Representationalism), or appeal to theoretical resources which are nowhere in sight, such as a well-grounded criterion for a genuinely “descriptive” statement. Finally, a viable defence of Representationalism would have to rest on an adequate theory of notions such as belief, assertion, commitment and judgement. But here, too, there are reasons to think that the most promising approaches will not respect Representationalism, being grounded on internal conceptions of representation, in the sense of §1.10, rather than external, world-tracking notions.

If there is a single key point, it goes something like this. In so far as our claims are representational, it seems plausible to assume that they are uniformly representational, whatever the subject-matter – in other words, that representation is a univocal notion, in this sense. But if representation is viewed as relation to our natural environment, univocity leads to the placement problem in an acute form. The problem is solved by abandoning the external notion of representation in favour of an internal notion; by recognising that the grip of the alternative picture rests in large part on the disquotational platitudes; and by insisting that we theorise about our relations to our natural environment in a different, non-semantic vocabulary. So long as we practice our naturalism in another key – in the pragmatic, functional dimension that opens up when we

abandon Representationalism – we retain univocity where it matters, while avoiding the placement problems altogether.

So, where do we stand, if we take this approach seriously? In one sense, as I said, we find ourselves among well-known landmarks. Close at hand are the Humean intuition that metaphysical puzzles tend to arise from mistakes about the genealogy of our beliefs and commitments; familiar forms of semantic and metaphysical deflationism; and the resources of an inferentialist account of assertion. Our destination turns out to be easily reached from any of these familiar viewpoints, for it is simply the region on the map where they all intersect. To find it, however, we had to remove the obstacle that has been hiding it from view. We had to move the mirror aside.²¹

References

- Abrams, M. H., 1953: *The Mirror and the Lamp: Romantic Theory and the Critical Tradition*, New York: Oxford University Press.
- Blackburn, S. 1984: *Spreading the Word: Groundings in the Philosophy of Language*, Oxford: Oxford University Press.
- Blackburn, S. 1993a: *Essays in Quasi-Realism*, New York: Oxford University Press.
- Blackburn, S. 1993b: ‘Realism, quasi, or queasy?’, in Haldane and Wright (1993), 365–383.
- Blackburn, S. 1998: *Ruling Passions: A Theory of Practical Reasoning*, Oxford: Oxford University Press.
- Brandom, R., 1994: *Making it Explicit*, Cambridge, Mass.: Harvard University Press.
- Brandom, R., 2000: *Articulating Reasons: An Introduction to Inferentialism*, Cambridge, Mass: Harvard University Press.
- Carnap, R., 1950: ‘Empiricism, semantics and ontology’, *Revue Internationale de Philosophie* 4, 20–40. Reprinted in *Meaning and Necessity: A Study in Semantics and Modal Logic*, 2nd enlarged edition, Chicago: University of Chicago Press, 1956, 205–221.
- Chalmers, D., 1996: *The Conscious Mind: In Search of a Fundamental Theory*, New York: Oxford University Press.
- Chomsky, N, 1957: *Syntactic Structures*, The Hague: Mouton.
- Chomsky, N, 1975: *Logical Structure of Linguistic Theory*, New York: Plenum; University of Chicago Press
- Chomsky, N, 1995: ‘Language and nature’, *Mind*, 104, 1–61.
- Haldane, J., and Wright, C. (eds.), 1993: *Reality, Representation, and Projection*, Oxford: Oxford University Press.

²¹I am grateful to Bob Brandom, Patrick Greenough, Jenann Ismael, Uriah Kriegel, Luca Moretti, Richard Rorty, Kevin Scharp, Lionel Shapiro and Michael Williams for comments on earlier versions of this chapter.

- Jackson, F., 1997: 'Naturalism and the fate of the M-worlds II', *Proceedings of the Aristotelian Society*, Supp. Vol. 71, 269–282.
- Johnston, M., 1993: 'Objectivity refigured: pragmatism without verificationism', in Haldane and Wright (1993), 8–130.
- Leiter, B., 2006: 'Introduction', in Brian Leiter, ed., *The Future for Philosophy*, Oxford: Oxford University Press, 1–23.
- Menzies, P. & Price, H., 1993: 'Causation as a secondary quality', *British Journal for the Philosophy of Science* 44, 187–203.
- Moore, G. E., 1903: *Principia Ethica*, Cambridge: Cambridge University Press.
- Price, H., 1988: *Facts and the Function of Truth*, Oxford: Basil Blackwell.
- Price, H., 1991: 'Agency and probabilistic causality', *British Journal for the Philosophy of Science* 42, 157–176.
- Price, H., 1998: 'Three norms of assertibility, or how the MOA became extinct', in Tomberlin, J., ed., *Philosophical Perspectives* 12, 41–54.
- Price, H., 2001: 'Causation in the special sciences: the case for pragmatism', in Domenico Costantini, Maria Carla Galavotti and Patrick Suppes, eds., *Stochastic Causality*, Stanford: CSLI Publications, 103–120.
- Price, H., 2007a: 'Quining naturalism', *Journal of Philosophy* 104, 375–405.
- Price, H., 2007b: 'Causal perspectivalism', in Huw Price and Richard Corry, eds., *Causation, Physics and the Constitution of Reality: Russell's Republic Revisited*, Oxford: Oxford University Press, 250–292.
- Price, H., 2008: *René Descartes Lectures*, Tilburg. Available online here: <http://philsci-archive.pitt.edu/archive/00004430/>
- Soames, S., 1989: 'Semantics and semantic competence', *Philosophical Perspectives*, 3, 575–596.
- Smith, B., 1992: 'Understanding Language', *Proceedings of the Aristotelian Society*, 109–141.
- Stich, S. 1996: *Deconstructing the Mind*, New York: Oxford University Press.
- Williams, B., 1973: 'Consistency and Realism', in *Problems of the Self*, Cambridge: Cambridge University Press, 187–206.
- Williams, M., 1999: 'Meaning and Deflationary Truth', *Journal of Philosophy* 96, no. 11, 545–564.
- Williamson, T., 2006: 'Past the Linguistic Turn?', in Brian Leiter, ed., *The Future for Philosophy*, Oxford: Oxford University Press, 106–128.
- Wilson, J. 1832: 'Tennyson's Poems', *Blackwood's Magazine*, 31, 721–741.
- Wright, C., 1992: *Truth and Objectivity*, Cambridge, Mass.: Harvard University Press.